HITACHI
Inspire the Next

WHITE PAPER

Optimize Hitachi Storage and Server Platforms in VMware vSphere 6.5 Environments

Best Practices Guide

By Tsuyoshi Inoue and Jeff Chen

September 2017

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision History

Revision	Changes	Date
AS-619-00	Initial release	September 8, 2017

Table of Contents

Optimize Hitachi Storage and Server Platforms in VMware vSphere 6.5 Environments

Best Practices Guide

Use these best practices to implement Hitachi storage or Hitachi Unified Compute Platform in a VMware ecosystem. The focus is on environments which leverage SAN-based datastores. If you use Hitachi NFS datastores, consult Hitachi NAS Platform Best Practices Guide for NFS with VMware vSphere (MK-92HNAS028-01 or later, PDF).

Hitachi is an Elite Partner in VMware's Technology Alliance Partner program, a participant in VMware Ready Partner programs for Storage Infrastructure Services, and an embedded OEM partner. Together, Hitachi and VMware are committed to providing innovativue, business-enabling, with end-to-end virtualization solutions for the data center.

Through Hitachi platforms, products and native VMware integrations, Hitachi provides a complete solution for rapidly expanding VMware environments.

These best practices provide guidance when you want are planning to implement or who have already implemented Hitachi products with VMware products. This aids in building a solution that helps to ensure a VMware environment that provides performance, scalability, reliability, usability, resilience, and recoverability when paired with Hitachi products.

Demanding economic and IT environments makes it important for you to provide high-performance, scalable, highly available, and easy-to-manage infrastructures to maximize business and IT efficiency. To do this, expand your use of virtualization to more business-critical and mission-critical (tier one and tier two) applications. This makes your deployed storage, servers, and software supporting this environment more important.

Hitachi industry-differentiated storage and server architecture are the ideal solution for VMware vSphere environments. The depth and variety of Hitachi storage and server products ensures that you can build scalable, reliable, and resilient server and VDI solutions that are easy to use.

These best practices cover some of the Hitachi storage and server products in Table 1, "Hitachi Data Systems Storage and Server Products," on page 2.

TABLE 1. HITACHI DATA SYSTEMS STORAGE AND SERVER PRODUCTS

Hardware	Product
Storage	These Hitachi Virtual Storage Platform (VSP) models: ■ Virtual Storage Platform G1500 ■ Virtual Storage Platform F1500 ■ Virtual Storage Platform G800 ■ Virtual Storage Platform G600 ■ Virtual Storage Platform G400 ■ Virtual Storage Platform G200 ■ Virtual Storage Platform F800 ■ Virtual Storage Platform F600 ■ Virtual Storage Platform F400 Hitachi NAS Platform (HNAS) Hitachi Accelerated Flash (HAF)
Servers	Hitachi Compute Blade 500 (CB 500) Hitachi Compute Blade 2500 (CB 2500) Rack Optimized Server for Solutions, 2U Four Node Rack Optimized Server for Solutions, 2U Single Node
Converged and Hyper-Converged	Hitachi Unified Compute Platform HC (UCP HC) Hitachi Unified Compute Platform 2000 (UCP 2000) Hitachi Unified Compute Platform 4000 (UCP 4000)

Some of the Hitachi software products covered by these best practices are listed in Table 2.

TABLE 2. HITACHI DATA SYSTEMS PLUGIN AND ADAPTER PRODUCTS

Software	Product
Hitachi Storage Software	Hitachi Storage Virtualization Operating System (SVOS): ▪ Hitachi Dynamic Provisioning (HDP) ▪ Hitachi Dynamic Tiering (HDT) ▪ Active flash Hitachi Thin Image (HTI) Hitachi ShadowImage Hitachi TrueCopy Hitachi Universal Replicator (HUR) Global-active device on Virtual Storage Platform Remote replication extended (for 3DC scenarios) Hitachi Command Suite (HCS) Hitachi Automation Director (HAD) Hitachi Unified Compute Platform Advisor Hitachi Unified Compute Platform Director (UCP Director)
Hitachi Adapters and Plugins for VMware	Hitachi Storage Provider for VMware vCenter Hitachi Storage Management Pack for VMware vRealize Operations Hitachi Storage Connector for VMware vRealize Orchestrator Hitachi Storage Content Pack for VMware vRealize Log Insight Hitachi Storage Plug-in for VMware vCenter Hitachi Storage Replication Adapter Hitachi Compute Connector for VMware vRealize Orchestrator Hitachi Compute Plug-in for VMware vCenter Hitachi Compute Content Pack for VMware vRealize Log Insight Hitachi Storage Adapter for VMware vSphere Storage API for Array Integration Hitachi Data Instance Director (HDID) with Hitachi Virtual Infrastructure Integrator

Solution Components

Use these Hitachi products for VMware deployment. See Products and Features for more information regarding to other Hitachi products.

Hitachi Unified Compute Platform

Hitachi Unified Compute Platform (UCP) is a family of integrated and flexible reference solutions. Each Unified Compute Platform solution, configured for immediate deployment, runs top tier infrastructure applications without over-purchasing or provisioning unnecessary equipment. The entire solution is stack certified and compatible.

Unified Compute Platform 4000 is the foundation for business transformation and IT modernization that can host any application, at any scale, at any location. It includes Hitachi Unified Compute Platform Director, which is natively integrated with VMware vCenter and Microsoft® System Center to centralize and automate the management of your server, storage and networking components. Unified Compute Platform 4000 offers these capabilities in an agile, resilient, and scalable converged infrastructure.

With Unified Compute Platform Director as its orchestration layer, you can configure up to eight Unified Compute Platform 4000 systems in VMware environments to build very large scale virtualized environments. This dramatically reduces the number of VMware vCenter licenses required.

You can meet a true zero recovery time objective (RTO) and recovery point objective (RPO) service level using Unified Compute Platform 4000 with global-active device supported in Unified Compute Platform Director.

Hitachi Unified Compute Platform Director

Hitachi Unified Compute Platform Director is a powerful single point of management tool offering end-to-end visibility of the entire converged virtualization infrastructure. Unified Compute Platform Director enables management of Hitachi Unified Compute Platform, providing a unified view of resources from a single pane of glass.

Unified Compute Platform Director delivers visibility and management capability so you can align IT with business processes using familiar tools. The solution monitors all the elements of Unified Compute Platform in a single view and provides a health status of all solution elements.

You can monitor virtual resources rapidly and efficiently across a single datacenter or multiple datacenters with Unified Compute Platform Director.

Hitachi Unified Compute Platform Advisor

Hitachi Unified Compute Platform Advisor (UCP Advisor) brings simplified IT administration to virtualized, converged, and hyperconverged systems from Hitachi. UCP Advisor supports guided life-cycle management to the server, network, and storage elements within supported Unified Compute Platform systems.

Provide operational workflows for end-to-end provisioning of VMware ESXi hosts, datastores and virtual LAN (VLAN) synchronization management with Unified Compute Platform Advisor.

When using Unified Compute Platform Advisor, manage single or multiple supported Unified Compute Platform systems or logical systems.

Figure 1 on page 5 shows the user interface of Unified Compute Platform Advisor within the VMware vSphere web client.

Figure 1

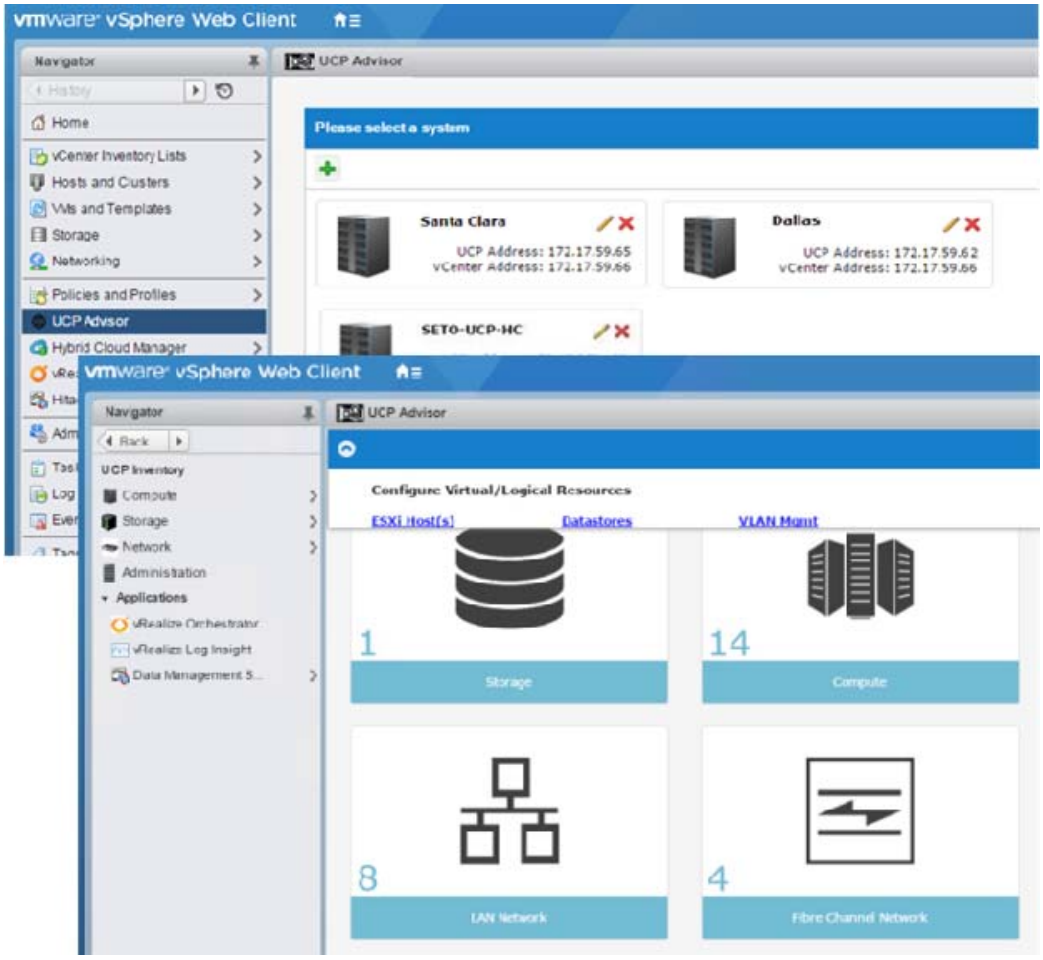

Hitachi Data Instance Director with Hitachi Virtual Infrastructure Integrator

Hitachi Data Instance Director (HDID) provides a modern, holistic approach to data protection, recovery, and retention.

Data Instance Director has a unique workflow-based policy engine, presented in a whiteboard-style user interface. It helps you to map copy data management processes to business priorities.

Data Instance Director includes a wide range of fully integrated storage-based and host-based incremental-forever data capture capabilities. These can be combined into complex workflows to automate and simplify copy data management.

When used with Hitachi Virtual Infrastructure Integrator, Data Instance Director provides effective data protection with the following:

- Virtual-machine-focused storage-snapshot-based backup
- Zero backup window
- Reduced storage consumption
- Support for block and file VMware deployments

Virtual Infrastructure Integrator and Data Instance Director simplifies operations to do the following:

- Integrate VMware user interface experience
- Flexibly backup groups and policies

Hitachi Automation Director

Hitachi Automation Director works with the Hitachi Command Suite management infrastructure, providing configurable best-practice-based service templates for simple, application-specific provisioning of storage resources to databases, applications, and VDI environments. Built-in support for role-based access control means that, once established, these templates can be used to move to self-service provisioning.

- Automate complex provisioning for enhanced efficiency with smart provisioning
- Take advantage of tiered storage resources and reduce costs
- Customize service templates with role-based privileges for IT administrators

Hitachi Adapters for VMware Environments

Hitachi adapters for VMware environments can centralize and automate management of all the hardware elements in VMware vSphere through a single pane of glass. Take advantage of intuitive performance information. Information on the health, risk and efficiency of your data center infrastructure simplifies maintenance and makes your team more proactive.

Hitachi Storage Provider for VMware vCenter

Use Hitachi Storage Provider for VMware vCenter to deploy VMware vSphere Virtual Volumes for policy-controlled software-defined datacenter. It provides granular provisioning and management of virtual machines. There are application-specific data services, such as snapshot, replication, and so forth. Reduce operational burden between virtual infrastructure administrator and storage administrator with this provider.

Hitachi Storage Management Pack for VMware vRealize Operations

Manage physical and virtual infrastructure using Hitachi Storage Management Pack for VMware vRealize Operations. Reduce mean time to resolution (MTTR) and operational costs. This tool provides deep visibility into Hitachi storage components through the following dashboards:

- Top consumers
- Performance
- Capacity

Hitachi Storage Connector for VMware vRealize Orchestrator

Automate and orchestrate tasks on Hitachi storage with Hitachi Storage Connector for VMware vRealize Orchestrator, supporting both block and file. Enhance the capabilities of VMware vRealize Orchestrator by providing access to Hitachi storage-specific workflows such as the following:

- Provision storage datastore or LUN
- Provision NFS datastores
- Back up SAN datastores

Hitachi Storage Content Pack for VMware vRealize Log Insight

Hitachi Storage Content Pack for VMware vRealize Log Insight deliver real-time log management for intelligent analysis and better troubleshooting across physical and virtual infrastructure. It simplifies searching for errors by collecting and grouping information to show important, relevant, and useful events. You are provided a comprehensive view into Hitachi storage systems, enabling them to spot potential issues and keep track of components showing departure from normal operation logs.

Hitachi Storage Plug-in for VMware vCenter

Using Hitachi Storage Plug-in for VMware vCenter, integrate management of Hitachi storage systems within the VMware vCenter console. This allows your VMware vCenter administrator to automatically configure VMware ESXi hosts to Hitachi storage systems. Use it to provide visibility into mapping of VMware vCenter datastores to Hitachi storage system logical drives.

Hitachi Storage Replication Adapter

With automated failover in disaster-recovery scenario, Hitachi Storage Replication Adapter also provides nondestructive disaster recovery testing.

Hitachi Compute Connector for VMware vRealize Orchestrator

Hitachi Compute Connector for VMware vRealize Orchestrator provides server configuration, discovery, and additional management functionality for customizing automation. Resolve operational issues quickly and reuse those solutions as best practices, thus avoiding errors and repetitive work.

Hitachi Compute Plug-in for VMware vCenter

Have deep visibility into Hitachi Compute Blade information in VMware vSphere Web Client using Hitachi Compute Plug-in for VMware vCenter. It provides detailed information on virtual machines running on Hitachi severs. There is the ability to perform power, logical identifier (LID), and remote console operations on Hitachi servers.

Hitachi Compute Content Pack for VMware vRealize Log Insight

Hitachi Compute Content Pack for VMware vRealize Log Insight provides real-time log management for VMware environments, with machine-learning-based intelligent grouping, high-performance search, and better troubleshooting. It collects and analyzes all types of Hitachi compute-systems-related data and provides intuitive dashboards, including overview, power, security, configuration and maintenance. Administrators have visibility to unauthorized logons, credential misuse, privilege escalations, and anomalies, to easily identify potential malicious activity.

VMware vSphere Storage APIs — Array Integration

VMware vSphere Storage APIs — Array Integration (VAAI) allow VMware vSphere environments to use advanced features of Hitachi storage arrays. Using vSphere Storage APIs provide a way to use those advanced storage capabilities from within the VMware interface. Processing is directly on the storage infrastructure.

These performance enhancements move the I/O load from the dependent VMware vCenter host platform into the storage controller. By offloading storage related operations off to the storage subsystem, it speeds up the datastore and VMDK provisioning operations. This frees virtualization management for more critical tasks.

When used with VMware vSphere 5.x and 6.x, Hitachi storage supports the following API primitives:

- **Full copy** — This primitive enables the storage system to make full copies of data within the storage system without having the VMware ESXi host read and write the data.
- **Block zeroing** — This primitive enables storage systems to zero out many blocks to speed provisioning of virtual machines.
- **Hardware-assisted locking** — This primitive provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large VMware ESXi host farms sharing a datastore.
- **Thin provisioning stun** — This primitive enables the storage system to notify the VMware ESXi host when thin provisioned volumes reach certain capacity utilization threshold. When enabled, this allows the ESXi host to take preventive measures to maintain virtual machine integrity.
- **UNMAP** — This primitive enables a VMware ESXi host to inform the Hitachi storage array that space can be reclaimed that previously had been occupied by a virtual machine that has been migrated to another datastore or deleted.

Best Practices

Begin SAN design with reliability and scalability. A well-designed SAN will be able to recover quickly in the event of a single device failure. Also, a well-designed SAN will grow easily, as the demands of the infrastructure that it serves increases.

General VMFS Provisioning Best Practices

These are best practices for general VMFS provisioning.

Host Group and Host Mode Options

To grant a host access to an LDEV, assign a logical unit number (LUN) within a host group. These are the settings and LUN mapping for host group configurations.

Fibre Channel Port Options

If connecting a Fibre Channel port using a SAN switch or director, you must change the following settings in Hitachi Storage Navigator:

- **Port security** — Set the port security to **Enable**. This allows multiple host groups on the Fibre Channel port.
- **Fabric** — Set fabric to **ON**. This allows connection to a Fibre Channel switch or director.
- **Connection Type** — Set the connection type to **P-to-P**. This allows a point-to-point connection to a Fibre Channel switch or director.

Hitachi recommends you apply the same configuration to a port in cluster 1 as to a port in cluster 2 in the same location. For example, if you create a host group for a host on port CL1-A, also create a host group for that host on port CL2-A.

One Host Group per VMware ESXi Host Configuration

If you plan to deploy VMware ESXi hosts, each host's WWN can be in its own host group. This approach provides granular control over LUN presentation to ESXi hosts. This is the best practice for SAN boot environments, because ESXi hosts do not have access to other ESXi hosts' boot LUNs.

However, in a cluster environment, this approach can be an administrative challenge because keeping track of which ESXi hosts are in a cluster can be difficult. When multiple ESXi hosts need to access the same LDEV for clustering purposes, the LDEV must be added to each host group. This operation is error prone and might lead to confusion.

If you have numerous ESXi hosts, this approach can be tedious.

One Host Group per Cluster, Cluster Host Configuration

VMware vSphere features such as vMotion, Distributed Resource Scheduler, High Availability, and Fault Tolerance require shared storage across the VMware ESXi hosts. Many of these features require that the same LUNs that are presented to all ESXi hosts participating in these cluster functions.

If you plan to use ESXi hosts with these features, create host groups with clustering in mind. Place all the WWNs for the clustered ESXi hosts in a single host group. This ensures that when adding LDEVs to the host group, all ESXi hosts see the same LUNs. This creates consistency with LUN presentation across all hosts.

Host Group Options

On the Hitachi Virtual Storage Platform family storage, create host groups using Hitachi Storage Navigator. Change the following host mode and host mode options to enable VMware vSphere Storage APIs — Array Integration (VAAI):

- **Host Mode** — 21[VMware Extension]
- **Host Mode Options**:
  - Enable 54-(VAAI) Support Option for the EXTENDED COPY command
  - Enable 63-(VAAI) Support option for vStorage APIs based on T10 standards

Change the following host mode and host mode options to enable a VMware Virtual Volumes environment on Hitachi storage using vSphere Storage APIs — Array Integration:

- **Host Mode** — 21[VMware Extension]
- Host Mode Options:
  - Enable 63-(VAAI) Support option for vStorage APIs based on T10 standards
  - Disable the VAAI rules on ESXi hosts. See Hitachi Storage Provider for VMware vCenter Deployment Guide on how to setup VAAI on ESXi hosts.

For more details regarding VMware Virtual Volume configurations, see VMware vSphere 6.5 Virtual Volume Environment Deployment in Hitachi Virtual Storage Platform.

Zoning

Use zoning to enable access control in a SAN environment. Through zoning, a SAN administrator can configure which HBA WWPNs on the VMware ESXi host can connect to which WWPNs on the Hitachi storage array storage processors.

Zoning Configurations

You can break zoning down into the following different configurations:

- **Single Initiator to Single Target zones** — This configuration allows one initiator to be zoned to only one target. This configuration is the most resilient configuration, as traffic originating from another Initiator on the SAN will have less impact than the initiator in this zone.
- **Single Initiator to Multiple Target Zones** — This configuration allows one initiator to be zoned to multiple targets in a single zone.
- **Multi Initiator Zones** — This configuration allows multiple initiators to be zoned to multiple targets in a single zone. This exposes all initiators and targets to all traffic in the zone. Events such as a malfunctioning HBA could affect all initiators and targets in the zone and either negatively affect performance for all or bring down the Fibre Channel network completely. **This configuration is never recommended.**

From lab testing at Hitachi on VMware environment using Hitachi Virtual Storage Platform family storage, there is no significant performance difference found in between single-initiator to single-target and single-initiator to multiple-targets. However, consider the following when you decide.

- Advantage of single-initiator to single-target
  - There is slightly easier Fibre Channel management with simpler zone configuration.
  - It is easier to troubleshoot a storage issue.
  - For some other operating systems, it might be an issue if you present the same LUN as different LUN number. (This does not apply to VMware.)
- Advantage of single-initiator to multi-targets
  - This provides extra redundancy in case of storage CHB (channel board) failure.
  - It distributes I/O more evenly to multiple storage ports in a mixed VMware environment.

Hitachi recommends using single initiator to single target as it is the most secure and stable, but single initiator to multiple target is also acceptable.

Mutipathing

Multipathing allows a VMware ESXi host to use more than one physical path between the ESXi host and the storage array.

Multipathing provides load balancing. This is the process of distributing I/O across multiple physical paths, to reduce or remove potential bottlenecks.

Multipathing also provides redundancy and fault tolerance in case of a failure of any element in the SAN network, such as an adapter, switch, or cable. The ESXi host can switch to another physical path not using the failed component. This process of path switching to avoid failed components is known as path failover.

To support path switching with a Fibre Channel SAN, the ESXi host typically has two or more HBAs available from which the storage array can be reached. It also has full fault tolerance that uses two or more switches.

Additionally, for full fault tolerance, two storage processors on Hitachi storage arrays should be utilized so that the HBA can use a different path to reach the disk array.

Path selection and multipathing policies supported by ESXi hosts are the following:

- **Most Recently Used (VMware)** — This policy uses the last successfully used path. If the last successfully used path is not available, then path failover occurs and the next available path is used. This new path continues to be used until it is no longer available, even if the previously used path becomes available again.
- **Round Robin (VMware)** — This policy sends a set number of I/O down the first available path, then sends the same set number of I/O down the next available path. This repeats through all available paths, and then starts over again and repeats. If a path becomes unavailable, it is skipped over until it becomes available again.
- **Fixed (VMware)** — This policy has a preferred path that can be set by the VMware vCenter administrator. This path is used until it becomes unavailable. Then, it fails over to the next available path until it becomes unavailable. In which case, the path fails over to the next available path, or until the preferred path becomes available again. If the preferred path does become available again, then the system fails back to the preferred path.
- **Hitachi Dynamic Link Manager** — VMware ESXi also supports third party path selection policies. Hitachi Dynamic Link Manager is Hitachi's multipathing software that integrates with global-active device on Hitachi Virtual Storage Platform and Hitachi High Availability Manager to provide load balancing and path failover capabilities for servers.

Hitachi recommends using the **round robin multipathing** policy. This multipathing policy takes advantage of all available paths and bandwidth. Taking advantage of all available paths assures maximum performance from the SAN infrastructure.

VMware ESXi Round Robin Path Selection Plug-in Path Rotation IOPS Limit

The VMware ESXi Round Robin Path Selection Plug-in (PSP) balances the load across all active storage paths. A path is selected and used until a specific quantity of I/O has been transferred.

The I/O quantity at which a path change is triggered is known as the limit. After reaching that I/O limit, the PSP selects the next path in the list.

The default I/O limit is 1000, but can be adjusted if needed to improve performance. Specifically, it can be adjusted to reduce latency seen by the ESXi host when the storage system is not seeing latency.

The recommended PSP limit is 1000. Setting the limit to 1 has not shown any improvement in IOPS or latency when testing with Hitachi storage systems.

Redundant HBA

In a VMware vSphere environment, the ESXi hosts should have two or more HBAs. Not only will this allow for greater I/O throughput on the SAN as more paths are available when using the round robin (VMware) multipathing policy, multiple HBAs also allow for redundancy and greater reliability in case of a component failure.

Redundant Fibre Channel Switches

When designing and building a reliable and scalable SAN environment, multiple Fibre Channel switches are recommended.

Each VMware ESXi host should have an equal number of connections to each Fibre Chanel switch. Each Hitachi storage array should have an equal number of connections from each storage processor to each switch.

The ESXi host port in the Fibre Channel HBA is referred to as the initiator. The storage processor port in the Hitachi storage array is referred to as the target.

This SAN Fibre Channel switch configuration ensures that a single switch failure will not leave an ESXi host unable to connect to a datastore, unable to continue running the virtual machines on those datastores.

It is recommended that the Fibre Channel switches not be up-linked to each other, creating separate Fibre Channel networks. This insures that conditions on a Fibre Channel network do not affect traffic on another Fibre Channel network, such as would happen with a malfunctioning HBA. This helps ensure great reliability.

Permanent Device Loss Autoremove

Permanent device loss (PDL) is a situation that can occur when a disk device either fails or is removed from the VMware vSphere host in an uncontrolled fashion.

Starting with VMware vSphere 5.5, an advanced setting called **AutoremoveOnPDL** was introduced and is implemented by default. This enables vSphere to remove devices marked as PDL to help prevent reaching the 255-device limit for an VMware ESXi host.

However, if the PDL scenario is resolved and the device returns, then the ESXi host's storage system must be rescanned before the device will appear.

For VMware vSphere Metro Storage Cluster configuration, VMware recommends disabling **Disk.AutoremoveOnPDL** for vSphere 5.5 hosts by setting the host advanced settings value to **0**. For vSphere 6.0 hosts, this advanced setting is no longer required to be changed from the default configuration to properly recover the devices marked as PDL.

Visit VMware vSphere Metro Storage Cluster Recommendation Practice (PDF) for more information.

Table 3 shows the recommended settings for **AutoremoveOnPDL**.

TABLE 3. RECOMMENDED SETTINGS FOR AUTOREMOVEONPDL

Setting	VMware vSphere 5.5	VMware vSphere 6.x
By Default	Enable	Enable
Non-vMSC	Enable or Disable No evidence or recommendation from VMware document.	Enable
vMSC	Disable	Enable

UNMAP

In VMware vSphere 6.5, ESXi supports manual and automatic asynchronous reclamation of free space on VMFS5 and VMFS6 datastores. It automatically issues the **UNMAP** command to release free storage space in background on thin-provisioned storage arrays that support UNMAP operations.

The reclamation of the free space of the datastore can still be executed manually or scripted with following command:

```
esxcli storage vmfs unmap -l <datastore-name>
```

Note — Hitachi intends to support automatic **UNMAP** soon.

iSCSI Provisioning

iSCSI initiators and targets use TCP to create relationships called “sessions.” The initiator sees one logical connection to the target. An iSCSI session might also contain multiple logical connections.

From a VMware vSphere host perspective, the sessions might also be thought of in term of paths between the initiator and target. Having multiple connections per session enables the aggregation of bandwidth and can also provide load balancing.

Although vSphere does not support multiple connections per session, by configuring multiple sessions, you can configure load balancing and ensure path redundancy. See VMware vSphere Storage Guide 6.5 (PDF) for more details.

Failover with Software iSCSI

With software iSCSI, you can use multiple NICs that provide failover and load balancing capabilities for iSCSI connections between your host and storage systems.

Multipathing plug-ins do not have direct access to physical NICs on your host. So, for this setup, you first need to connect each physical NIC to a separate VMkernel port. You then associate all VMkernel ports with the software iSCSI initiator using a port binding technique. As a result, each VMkernel port connected to a separate NIC becomes a different path that the iSCSI storage stack and its storage-aware multipathing plug-ins can use.

LUN Size

Prior to Hitachi Virtual Storage Platform Gx00 models and Virtual Storage Platform Fx00 models, the recommended LUN size was limited to 4 TB because of the storage remote replication capability. With VSP Gx00 models and VSP Fx00 models, this limitation has been removed. The following lists the current maximum LUN size:

- The maximum LUN size for VMware vSphere 6.x is 64 TB.
- The maximum LUN size for VSP Gx00 models and VSP Fx00 models is 256 TB with replication.
  - The LUN must to be within the dynamic provisioning pool.

VMware vSphere Storage APIs — Array Integration Atomic Test and Set (ATS)

A change in the VMFS heartbeat update method was introduced in VMware ESXi 5.5 Update 2 (VMFS5), to help optimize the VMFS heartbeat process. The legacy method involved plain SCSI reads and writes with the VMware ESXi kernel handling validation. The new method offloads the validation step to the storage system. This is like other offloads related to VMware vSphere Storage APIs — Array Integration (VAAI).

This optimization results in a significant increase in the volume of ATS commands the ESXi kernel issues to the storage system and resulting increased load on the storage system. Under certain circumstances, VMFS heartbeat using ATS may fail with false ATS miscompare events. This causes the ESXi kernel to verify again its access to VMFS datastores. This leads to “Lost access to datastore” messages.

The improvement of this issue is implemented in VMFS6. The following setting is recommended for a VMware vSphere 6.5 environment.

- Set ATS heartbeat **OFF** for vSphere 6.5 with VMFS5.
- Keep ATS heartbeat **ON** for vSphere 6.5 with VMFS6.

Refer to ESXi host loses connectivity to a VMFS3 and VMFS5 datastore (2113956) for more details and how to turn off ATS.

Storage Resource Management

VMware vSphere provides several features to resolve the virtual machine performance and capacity problems and help your administrative tasks around storage, such as vSphere Storage DRS, vSphere Storage I/O control, and vRealize Operations Manager. Hitachi storage also provides some storage-aware functionalities, such as Hitachi Dynamic Tiering and active flash to address similar issues. It is important to grasp the differences and what vSphere and Hitachi storage each resolves with its functionality.

Hitachi Dynamic Tiering and Active Flash

Using Hitachi Dynamic Tiering, you can configure a storage system with multiple storage tiers using different types of data drives. This includes the following:

- FMD
- SSD
- SAS
- SATA
- External volumes

This helps improve the speed, capacity, and cost of performance. Dynamic Tiering improves underline storage resources with following conditions.

- **Trigger** — Monitor the I/O load per page and relocate the page to the optimal tier.
- **Time** — Define a user-specified period of at least 30 minutes.
  - Real-time with the active flash feature
- **Granularity** — Relocate the storage tier with a page size of 42 MB.

In addition, active flash monitors a page's access frequency level in real time to promote pages that suddenly became busy from a slower media to high-performance flash media.

In a VMware environment, many workloads tend to be highly random with smaller block size. This may not be suitable for deduplication and compression, even with all flash configuration. Hitachi Dynamic Tiering with active flash may a good option to improve capacity and cost by efficiently using the flash tier minimally.

VMware vSphere Storage DRS

A datastore cluster is a collection of datastores with shared resources and a shared management interface. When you create a datastore cluster, you can use VMware vSphere Storage DRS to manage storage resources.

Storage DRS generates recommendations or performs Storage vMotion migrations to balance space use across the datastore cluster. It also distributes I/O within the datastore cluster and helps alleviate high I/O load on certain datastores.

- **Trigger** — This happens when space use on a datastore exceeds a threshold, or I/O latency on a datastore exceeds a threshold.
- **Time** — Storage DRS is invoked at the configured frequency which is, by default, every eight hours. Or, it is invoked when one or more datastores in a datastore cluster exceeds the user-configurable space utilization thresholds.
- **Granularity** — VMware vCenter Server uses Storage vMotion to migrate virtual machine disks to other datastores in the datastore cluster to balance the resources.

When deciding which datastores to group into a datastore cluster, try to choose datastores that are as homogeneous as possible in terms of the following:

- Host interface protocol, such as FCP, iSCSI, NFS
- RAID level
- Performance characteristics

VMware recommends not mixing SSD and hard disks in the same datastore cluster. However, this does not apply to the datastores provisioned from a Hitachi Dynamic Tiering pool.

The following are recommendations for VMware vSphere Storage DRS with Hitachi storage:

- Enable only **Space** metrics when a datastore cluster contains multiple datastores that are provisioned the same dynamic provisioning pool with or without Hitachi Dynamic Tiering.
  - Moving a noisy neighbor within the same dynamic provisioning pool does not improve the performance.
- Enable **Space** and **I/O** metrics when a datastore cluster contains multiple datastores that are provisioned from different dynamic provisioning pools.
  - Moving a noisy neighbor to the other dynamic provisioning pool balances out the performance.

VMware vSphere Storage I/O Control (SIOC)

VMware vSphere Storage I/O Control (SIOC) extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation.

- **Trigger (Time)** — This happens when device latency exceeds a threshold.
- **Granularity** — Each virtual machine (or virtual disk) that accesses that datastore is allocated I/O resources in proportion to their shares.

In a mixed VMware environment, increasing the HBA LUN queue depth will not solve a storage I/O performance issue. It may overload the storage processors on your storage systems. The best practice from Hitachi Data Systems is to set the default HBA LUN queue depth to 32.

VMware vRealize Operations

VMware vRealize Operations helps visualize all the resources associated with a virtual machine in a single plane of grass. It bridges the gaps between virtual and physical object to identify where the problem takes place. The following are some characteristics of storage management with vRealize Operations:

- **Trigger** — This is when a threshold set in each monitoring item is exceeded.
- **Time** — This is the time to refresh and retrieve the information from storage. A five minutes interval is the default.
- **Granularity** — Set monitoring items and threshold values on a per storage object basis.

Hitachi Storage Management Pack for VMware vRealize Operations extends the functionality of vRealize Operations by providing simplified management capabilities for Hitachi storage components, improving performance and operational efficiency. It provides better visibility into performance and storage capacity planning for your end-to-end virtual infrastructure environment. Refer to the Hitachi Storage Management Pack for VMware vRealize Operations datasheet for more information.

Capacity Savings with the All-Flash Hitachi Virtual Storage Platform Fx00 Models and Hybrid Virtual Storage Platform Gx00 Models

Hitachi Virtual Storage Platform Fx00 and hybrid Virtual Storage Platform Gx00 models deliver superior all-flash performance for business-critical applications. It guarantees continuous data availability with a combination of high-density flash module drives (FMD) DC2. These drives use patented flash I/O management and specialized offload engines to maximize flash utilization.

The key factor affecting accommodation on a flash device is not performance, but capacity. So, this makes key factors the high raw capacity which the flash device has and the saving ratio which comes from deduplication and compression functionalities.

Capacity Saving with Deduplication and Compression Options

Regarding deduplication and compression, the Hitachi Virtual Storage Platform family has mainly two types;

- Hitachi Storage Virtualization Operation System (SVOS) provides software-based deduplication and/or compression with post processing
- FMD DC2 hardware based compression with Inline processing
- When you use FMD DC2 hardware-based compression, enabling the accelerated compression option on all parity groups of FMD DC2 drives is required.
- You can use either software based or hardware based deduplication and compression, or the combination of both. With the combination of both options, software based deduplication and hardware based compression are used.

Compression Recommendation and Considerations

Regarding compression option, using FMD DC2 hardware based compression is recommended for the following reasons:

- No performance degradation appears due to the truly hardware-offloaded in-line or real-time accelerated compression.
- Regarding the compression saving ratio, the differences between software-based and hardware-based are insignificant.
- Inline processing-based compression provides you with reduction of initial capacity and cost. You can estimate the required FMD DC2 capacity with the Hitachi Data Reduction Estimator. (Read more about using this tool, including how to get access to it.)

Software-based compression consumes extra storage compute resources. This post processing-based compression requires full allocated capacity to temporarily store for the initial phase as well.

Deduplication Recommendation and Considerations

Deduplication is highly effective in the virtualization environment, which tends to have duplicated data. This includes data such as the same operating system images, templates, and backups.

From lab validation results at Hitachi, enabling deduplication achieved a 60-70% capacity saving for a datastore where 8 virtual machines with an operating system VMDK resides (Microsoft® Windows Server® 2012 R2).

Enabling FMD DC2 hardware accelerated compression enhances deduplication with more than a 20% capacity saving. This combination of deduplication and compression achieved more than 80-90% capacity savings in total.

You can also estimate saving ratio and deduped capacity with the Hitachi Data Reduction Estimator. (Read more about using this tool, including how to get access to it.)

A main concern related to deduplication is performance degradation. This comes from mainly the following two factors:

- It consumes extra storage compute resources to perform deduplication and metadata management.
- The garbage collection running as a background task also require processing overhead. This task may increase storage CPU (MP) usage from 2% to 15%.

The following are some of the considerations with regards to software-based deduplication:

- It may impact I/O performance. Verify the performance by utilizing best practices or the cache optimization tool (COT) tool before using the capacity saving function.
- Because approximately 10% of the capacity is used for metadata and garbage data, capacity saving function should be applied only when the saving is expected to be 20% or higher.
- In deduplication and compression, processing is performed per 8 KB. Therefore, if the block size of the file system is an integral multiple of 8 KB, then the capacity saving is likely to be effective.
- The capacity saving function is not a good fit for high-write workloads. If the write workload rate is higher than garbage collection throughput, then the storage cache write-pending increases, causing performance degradation.

The capacity saving effect vary depends on your application and workload. You need to know your application workload and suitability before enabling a capacity saving feature. Table 4 lists the possible use cases for the capacity savings.

TABLE 4. DEDUPLICATION CONSIDERATION FOR GENERAL USE CASES

Use Case	Description
Microsoft Office®	Because there are many identical file copies, deduplication is effective.
VDI	Deduplication is very effective because of operation system area cloning.
Database (TPC-H)	Deduplication is not effective because the database has unique information for each block.
Database (TPC-C)	For a database that has many data updates, garbage data is increased, so it is not suitable.
Image/video	Compressed by application.
Backup/archive	Deduplication is effective between backups.

Flash Module Drive DC2 Configurations and Recommendations

As mentioned in “Deduplication Recommendation and Considerations” on page 16, the key factor affecting accommodation on a flash device is not performance, but capacity. The required flash memory drive (FMD) DC2 capacity can vary, whether there is dedupeable or compressible data. You can estimate it by using the Hitachi Data Reduction Estimator. (Read more about using this tool, including how to get access to it.).

The following are some recommendations for FMD DC2:

- If your application requires high IOPS and low latency, and if your data is compressible, FMD DC2 accelerated compression (without dedupe) might be an option.
- RAID-6 is the recommended RAID level for pool-VOLs, especially for a pool where recovery time from a pool failure due to a drive failure is not acceptable.
- Configure a parity group across the drive-boxes to maximize the performance by increasing the number of back-end path.

VMware VASA Provider for VMFS and Virtual Volumes

VMware vSphere Virtual Volumes is based on an integration and management framework between VMware vSphere and the storage system introduced by vSphere 6.0. The incompatibility between vSphere and the array virtual machine management is the difference between the primary units for both sides, such as virtual disks (VMDK files) for the virtual machines and LUNs for the array. This can require sacrificing SLAs for virtual machines.

For example, virtual machines that reside on the same data store (LUN) are forced to use the same RPO-leveraging storage replication, despite different RPOs required for each virtual machine.

With Virtual Volumes, the virtual disk becomes the primary unit of data management at the storage system level. It becomes possible to execute storage operations with granularity and to provision native storage-systems-based data services to individual virtual machines.

VMware vSphere APIs for Storage Awareness

VMware vSphere APIs for Storage Awareness (VASA) is either supplied by third-party vendors or offered by VMware. It enables communication between VMware vCenter Server and underlying storage. Through VASA, the storage entities can inform vCenter Server about their configurations, capabilities, and storage health and events.

In return, in certain environments, VASA can deliver virtual machine storage requirements from vCenter Server to a storage entity and ensure that the storage layer meets the requirements.

Differences Between VASA 1.0 and VASA 2.0

Table 5 shows the differences between VASA 2.0 and VASA 1.0.

TABLE 5. DIFFERENCE BETWEEN VASA 2.0 AND VASA 1.0

VASA Version	Storage service Integration with VMware vCenter Server	Datastore	Component	Notes
2.0	■ Tag-based storage policy	VMFS	VASA provider (VP)	More extensible and flexible compared to VASA 1.0. Available with Hitachi VP 3.4 or later.
2.0	■ Advertisement of storage capabilities and data services that storage arrays offer. ■ Bidirectional communications between VMware ESXi and VMware vCenter Server on one side, and storage arrays. ■ Offload virtual machine operations from vCenter or ESXi to storage arrays such as power on or off, cloning, VMware snapshot. ■ Align with SPBM (Note 2)	VVol	VASA provider provided by OVF (Note 1)	Because of tightly integration with vCenter and storage arrays in terms of storage management perspective and a part of host level based virtual machine operations are offloaded though VASA on out-of-band communication, the high availability for out-of-band is important (Note 3).
1.0	■ Advertisement of storage capabilities and data services that storage arrays offer. ■ Deliver a one-way communication with fixed and static information.	VMFS/NFS		Less integration with vCenter and storage array than VASA 2.0 in terms of VM host level based operations.

Note 1: From version 3.4.0, VASA 1.0-based function is integrated in VASA 2.0-based provider and now Hitachi provides and delivers VASA 1.0 and VASA 2.0-based functionality in a single provider as an OVF.

Note 2: For SPBM, Hitachi storage capabilities defined on array-side and advertised by VASA scheme-managed storage capabilities for SPBM on Hitachi storage is terminology that does not impose the same systematic solutions. The implementation of SPBM differs, based on the administrators, datacenters, or sites where it is used.

Storage administrators and virtual storage administrators work together to define the various capabilities for their system prior to implementation.

When working with managed storage capabilities, the values storage administrators manually attribute are based on how the administrators define the value of each storage container and respective storage pool capabilities. Auto-generated capabilities result from the storage system design for SPBM.

From version 3.4.0, Hitachi storage capabilities can be defined and set in Hitachi Storage Provider for VMware vCenter using its unified web user interface. Earlier than 3.4.0, another user interface (Hitachi Command Suite-based) is required to define and set these capabilities.

Note 3: For high availability considerations for Virtual Volumes out-of-band, refer to “Hitachi Implementation for VASA 1.0 and VASA2.0” on page 20.

Note 4: For a tag-based storage policy, refer to “Tag-based Storage Policy (SPBM for VMFS)” on page 20.

Hitachi Implementation for VASA 1.0 and VASA2.0

Hitachi Storage Provider for VMware vCenter is Hitachi’s implementation of the VASA provider. From Hitachi Storage Provider for VMware vCenter version 3.4.0, VASA 1.0 0-based functionality is integrated in VASA 2.0-based provider. Now Hitachi provides and delivers both VASA 1.0-based and VASA 2.0-based functionality in a single provider as an OVF.

Hitachi Storage Provider for VMware vCenter supports the following functions in terms of management, configuration flexibility and high availability.

- Log on to Hitachi Storage Provider for VMware vCenter with VMware vCenter Server Single Sign On account.
- Register Hitachi Storage Provider for VMware vCenter to multiple vCenter Servers.
- Manage Hitachi Storage Provider for VMware vCenter with VMware vSphere Fault Tolerance or vSphere High Availability functions.
- With the VASA provider, enable monitoring of the application-level under the VMware vSphere High Availability configuration. By enabling the monitoring of the virtual machine and application option of vSphere High Availability, the virtual machine automatically restarts when the VASA provider service stops.

Tag-based Storage Policy (SPBM for VMFS)

VMware VASA 1.0 has less capability than VASA 2.0, delivering a one-way communication with fixed and static information. In a SPBM scheme, to advertise storage capabilities delivered from VASA 1.0, there is no automatic way to align with or interact with creation of the virtual machine storage policy. Leveraging PowerCLI-based tagging in VMware vSphere 6.x, Hitachi provides a differentiated function called tag-based storage policy for VMFS.

The tag-based storage policy provides the same outcome for VMFS datastores as storage policy based management (SPBM) does for VMware Virtual Volumes. You can set a storage capability profile on the pool that is serving the VMFS datastores or you can customize the storage capability profile of an individual LUN. Hitachi Virtual Storage Platform automatically tags the datastores in VMware vCenter for existing and new datastores. In a way like SPBM for VMware Virtual Volumes, you can create a virtual machine storage policy using tags with SPBM.

Hitachi Storage Capabilities Defined on Array-side and Advertised by VASA Scheme

Storage administrators and virtual storage administrators work together to define the various capabilities for their system prior to implementation.

From Hitachi Storage Provider for VMware vCenter v3.4.0, you can define and set Hitachi storage capabilities using Hitachi Storage Provider for VMware vCenter with its web user interface. Earlier than version 3.4.0, the user interface from Hitachi Command Suite is required to define and set these capabilities.

SPBM is a structure that is based on a defined virtual machine storage policy. The VASA provider from Hitachi automatically selects the storage resources that meet the virtual machine storage policy and creates virtual volumes.

Figure 2 on page 21 shows an example of a storage capability profile definition by VASA provider in the web user interface. Figure 3 on page 21 shows an example virtual machine storage policy definition by VMware vSphere Web Client.

Figure 2

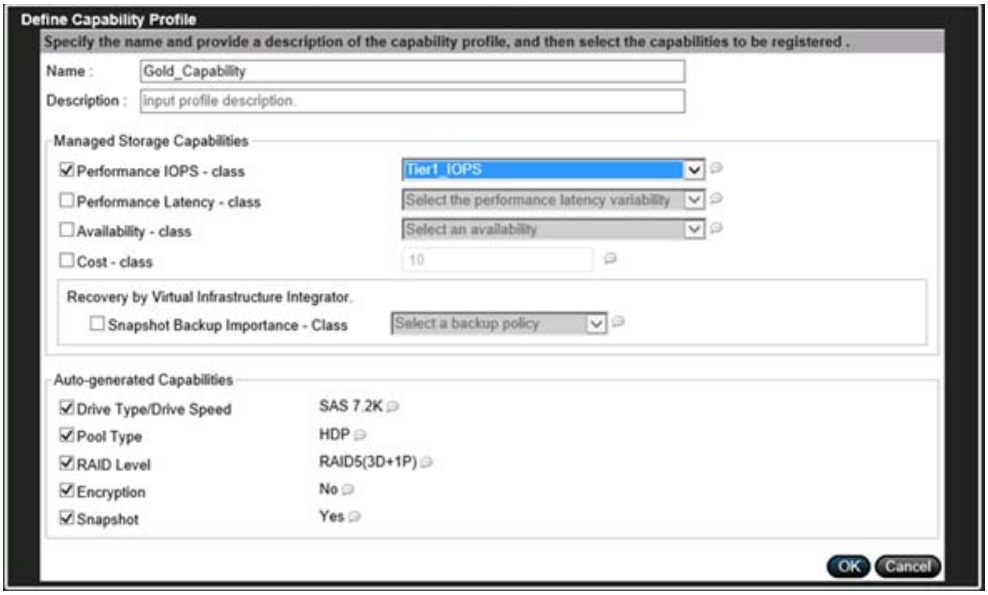

Figure 3

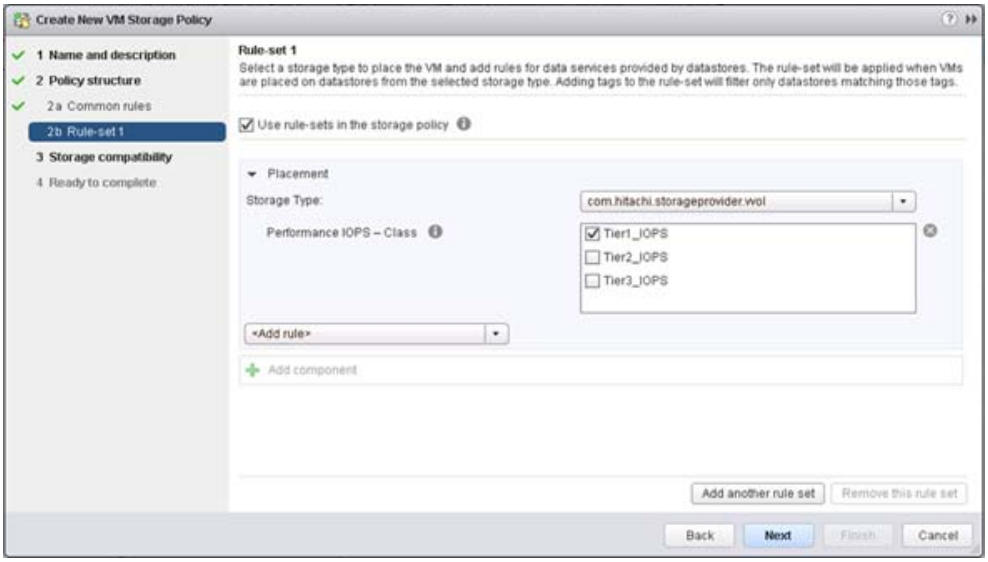

Recommendations to Support a VMware Virtual Volume Architecture

Do the following when deploying Hitachi Storage Provider for VMware vCenter to support a VMware Virtual Volume architecture:

- Deploy Storage Provider for VMware vCenter in a management cluster as a virtual machine.
- Deploy Storage Provider for VMware vCenter on the host using high-performance processor and memory.

Ensuring high availability for the following are key in the out-of-band design of a block environment for VMware Virtual Volumes from an operational standpoint:

- Hitachi Storage Provider for VMware vCenter
- The SVP required for Hitachi Virtual Storage Platform Gx00 models and Virtual Storage Platform Fx00 models.
- When one of the components above become unavailable, the only storage management operations related to virtual volume object metadata manipulation will be impacted, such as clone virtual machines, snapshot operations, and power operations.
  - The out-of-band communication on this Hitachi block implementation of VMware Virtual Volumes in a down situation does not affect virtual machine I/O, as the I/O flows through the data path (PEs).
- For Hitachi Storage Provider for VMware vCenter that is packed in an open virtual appliance (OVA) and deployed as a single virtual machine, the following are supported from Hitachi Storage Provider for VMware vCenter v3.3 or later:
  - VMware vSphere Fault Tolerance
  - VMware vSphere High Availability

  The recommendation is to enable vSphere Fault Tolerance or vSphere High Availability with Storage Provider for VMware vCenter deployed as a single virtual machine for continuous VMware Virtual Volume operations.
- Deploying the external SVP required for Hitachi Virtual Storage Platform Gx00 models and Virtual Storage Platform Fx00 models in a management cluster as a virtual machine is supported starting with storage microcode version 83-03-23/00. Also, enabling VMware vSphere Fault Tolerance with the external SVP virtual machine is supported and recommended for continuous VMware Virtual Volume operations. Enabling VMware vSphere High Availability with the external SVP will be supported in a later release.
- The VMware VASA provider enables monitoring of the application level under the VMware vSphere High Availability configuration. By enabling the monitoring of the virtual machine and application option of VMware vSphere High Availability, the virtual machine will automatically restart when the VASA provider service stops.

Multiple VMware vCenter Server Support

Hitachi Storage Provider for VMware vCenter can be registered with multiple VMware vCenter servers. This improves the flexibility of configuration. For example, the ability to have multiple vCenter server for production, test, or development environments respectively with a single shared storage array.

You can also register multiple VMware VASA providers in the VMware vCenter server. However, registering the same storage system in multiple VASA providers is not supported. A tag-based storage policy can be used with only one vCenter server.

Refer to the following documents for more details about virtual volumes with Hitachi storage:

- Hitachi Storage Provider for VMware vCenter Deployment Guide 3.4.0 (MK-90ADPTR010-18)
- VMware vSphere 6.5 Virtual Volume Environment Deployment in Hitachi Virtual Storage Platform (PDF)

VMware Site Recovery Manager Best Practices

These are the best practices for VMware Site Recovery Manager (SRM).

Standard Storage SRM and Stretched Storage SRM with Global-active Device Best Practices

VMware vCenter Site Recovery Manager integrates tightly with Hitachi storage arrays using Hitachi Storage Replication Adapter. This use provides centralized management of recovery plans. Tight integration between storage arrays, VMware vCenter, VMware vCenter Site Recovery Manager, and Hitachi Storage Replication Adapter ensure a coordinated recovery for large, business critical environments.

Remote data replication is a key function in building out stable and reliable disaster recovery environments. Replicating data to a remote secondary site represents the most effective insurance policy against a catastrophic failure. Although you can perform data replication at the server level, you can perform data replication more effectively within the storage infrastructure.

The following are two types of underlying storage defined by VMware SRM use and corresponded remote replication types which Hitachi supports through Hitachi Storage Replication Adaptor.

- Standard storage (active-standby solution): leveraging Hitachi TrueCopy or Hitachi Universal Replicator
- Stretched storage (active-active solution): leveraging global-active device on Hitachi Virtual Storage Platform

Table 6 lists the differences between two types of storage.

TABLE 6. COMPARISON BETWEEN STANDARD STORAGE AND STRETCHED STORAGE

Type	SRM with Standard Storage	SRM with Stretched Storage
Business continuity	Site failover is required with down time even though planned migration such as site maintenance is being conducted.	During planned migration such as site maintenance, no disruption and down time occurs by using Cross-vCenter vMotion with Stretched storage is made up by global-active device and Hitachi Storage Replication Adapter.
Storage availability	Site failover is required due to primary storage failure. It costs application downtime.	When primary storage failure occurs, no site failover is required by using the cross path to remote site storage which is virtualized as a single stretched storage and volume across the sites powered by global-active device technology.
Simplicity	Simple as traditional disaster recovery configuration consists of primary storage and secondary storage.	In addition to traditional disaster recovery configuration, there is a need to consider quorum storage and additional path between sites as cross-path, and so forth. It tends to be more complex and larger one.

You may decide, depending on required RPO, which results in which replication type you choose.

For more information, see Hitachi Unified Compute Platform 2000 for VMware vSphere with Site Recovery Manager solution using Stretched Storage (Global-Active Device) Reference Architecture Guide

VMware vSphere Metro Storage Cluster Best Practices

A VMware vSphere Metro Storage Cluster architecture on Hitachi storage platforms provides an ideal solution for maximizing availability and uptime by clustering physical datacenters within metro distances. The metro storage cluster solution from Hitachi Data Systems consists of storage systems presenting replicated storage as a single LUN from different geographically distributed sites. This design enables high availability of services by allowing virtual machine migration between sites with no downtime.

Changes in Multipathing and Path Configuration Best Practice

These are changes in multipathing and patch configuration best practices a VMware vSphere Metro Storage Cluster environment.

Resolution for PDL and APD

There have been cases which require manual host shut-down to invoke VMware vSphere High Availability fail over due to the lack of interaction between vSphere High Availability and permanent device lost (PDL) or all path down (APD). This is a key requirement for a massively high availability solution, such as VMware vSphere Metro Storage Cluster.

VMware vSphere 6.0 and later has a powerful new feature as part of vSphere High Availability called virtual machine component protection (VMCP).

VMCP protects virtual machines from storage related events, specifically permanent device loss (PDL) and all paths down (APD) incidents. In cases of PDL or ADP, a virtual machine is automatically failed over to another host in a high availability cluster by vSphere High Availability.

Global-active Device with Native Multi-Pathing with ALUA

For a VMware Metro Storage Cluster configuration, global-active device with VMware native multi-pathing (NMP) with ALUA is supported with micro code 83-03-01-x0/00 and later. This feature allows you to present I/O to the remote site storage across the long distances path which causes high response time by specifying it as non-optimized path. This minimizes response time and the cost of WAN traffic. It is recommended to turn on this feature with site distances greater than 20 miles (32 km).

Here is an example of enabling ALUA mode and specify non-optimized path on Hitachi storage

```
raidcom modify ldev -ldev_id XXXXXX -alua enable

raidcom modify lun -porrt CLX-C-X -lun_id all -asymmetric_access_state non_optimaized
```

Here is an example of enabling SATP rule set as ALUA for Hitachi devices and selecting PSP as round robin on VMware ESXi side.

```
esxcli storage nmp satp rule add -V HITACHI -M "OPEN-V" -P VMW_PSP_RR -s VMW_SATP_ALUA
-c tpgs_on
```

- Hitachi support recommends using RR instead of MRU.

Quorum Disk Enhancemnt

For pairs created, resynchronized, or swap resynchronized, if a failure occurs in the quorum disk, the pair status of the PVOL and S-VOL does not change from PAIR (Mirror RL) for the following:

- Microcode version 80-05-0x or later for Hitachi Virtual Storage Platform G1000, Virtual Storage Platform G1500, and Virtual Storage Platform F1500
- Firmware version 83-04-0x or later for Virtual Storage Platform Gx00 models
- Firmware version 83-04-2x or later for Virtual Storage Platform Fx00 models

In this stage, there is no impact on virtual machine while I/O continues to be processed on local storage on each site.

After the quorum disk is blocked, the pair is suspended when the remote paths are disconnected.

- The P-VOL status and the I/O mode change to PSUE (Local).
- The S-VOL status and the I/O mode change to PAIR (Block).
- The server I/Os continue in the P-VOL.

The pair might be suspended and the status and the I/O mode of the P-VOL and the S-VOL might change to PSUE (block) depending on the timing of the remote path disconnection after the quorum disk is blocked.

Uniform and Non-Uniform Host Access

While the Hitachi Storage Cluster for VMware vSphere solution supports uniform and non-uniform host access topology, Hitachi recommends uniform host access deployment where feasible for upmost high availability requirements.

- **Uniform host access configuration** — This is when VMware ESXi hosts from both sites are all connected to a storage node in the storage cluster across all sites. Paths presented to ESXi hosts are stretched across distance.
- **Non-uniform host access configuration** — This is when VMware ESXi hosts in each site are connected only to storage node or nodes in the same site. Paths presented to ESXi hosts from storage nodes are limited to the local site.

Refer to Implement vSphere Metro Storage Cluster using Hitachi Storage Cluster for VMware vSphere, featuring Hitachi Virtual Storage Platform G1000/G1500/F1500/Gx00/Fx00 (2145375) for more information regarding to this topic.

3DC with VMware Site Recovery Manager Best Practices

The 3DC solution consists of clustered primary and secondary datacenters leveraging global-active device in Hitachi Virtual Storage Platform, and the third data center which is replicated from the others as a disaster recovery site leveraging Hitachi Universal Replicator with delta resync functionality.

Hitachi Universal Replicator with delta resync functionality establishes storage remote replication from the primary data center to the third data center and from the secondary data center to the third data center respectively. This is called the global-active device 3DC delta resync environment, as shown in Figure 4 on page 26.

Figure 4

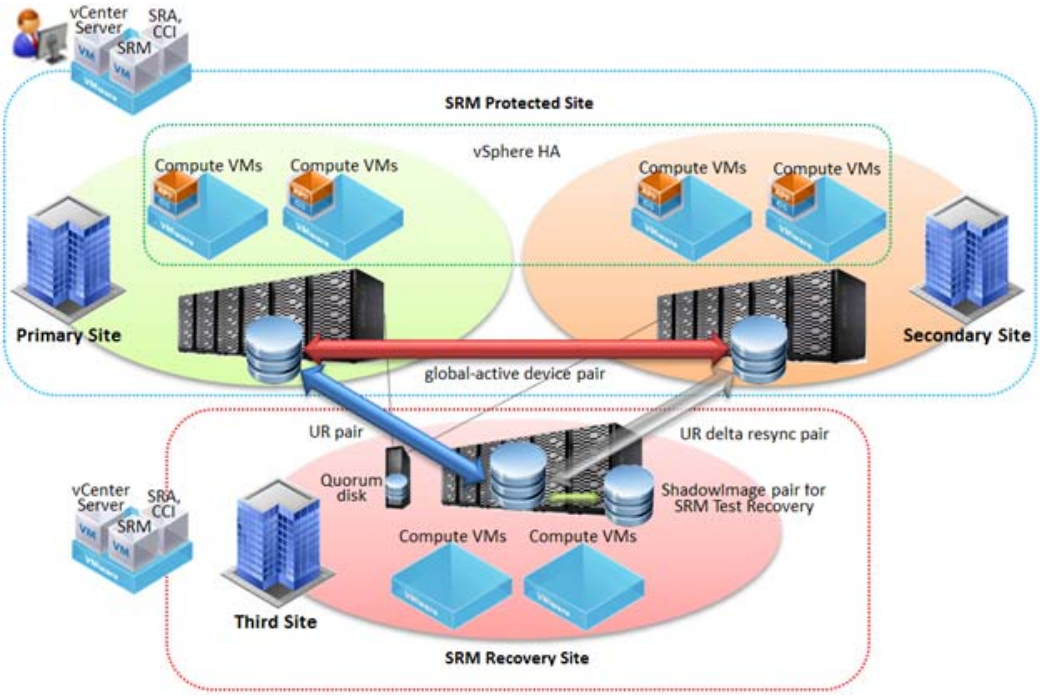

Installing VMware Site Recovery Manager in this 3DC environment gives you the orchestrated and repeatable planned or un-planned migration or disaster recovery operations using tested and proven recovery plan. This enables end-to-end virtual machine protection across 3 datacenters. As a normal state, VMware SRM protects the virtual machine between the primary and the third data center.

This solution is based on VMware Metro Storage Cluster, which clusters the primary and the secondary data centers within a single VMware vCenter data center object and using stretched storage cluster powered by global-active device.

When the primary datacenter goes down, the virtual machine can be restarted on the secondary data center, leveraging VMware vSphere High Availability fail over functionality as a VMware Metro Storage configuration. During failover from the primary to secondary datacenter, storage remote replication established from the primary to the third data center is also automatically failover to the other one established from the secondary to the third data center by leveraging delta resync functionality

For global-active device 3DC delta resync environment solution, the virtual machine protected by VMware SRM can follow this datacenter fail over movement and re-protected the virtual machine between secondary to third data center with minimizing effort.

Note — As a normal state, VMware SRM protects the virtual machine between primary and third data center. When the primary data center goes down, storage remote replication can automatically failover though, the re-protection of virtual machine by SRM requires some manual operation to switch source and target data center. To do this, switching command control interface configuration file and restarting the Hitachi Open Remote Control Manager daemon is required.

Products and Features

This describes the other hardware and software components mentioned in this best practice guide.

Hitachi Virtual Storage Platform

Hitachi Virtual Storage Platform Gx00 models are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Virtual Storage Platform Fx00 models deliver superior all-flash performance for business-critical applications, with continuous data availability. High-performance network attached storage with nondisruptive deduplication reduces the required storage capacity by up to 90% with the power to handle large, mixed-workload environments.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 and Virtual Storage Platform Fx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort.

To support read/write copies of the same data in two places at the same time, Virtual Storage Platform has global-active device. Its active-active design implements cross-mirrored storage volumes between two matched VSP systems that accept read/write I/Os on both sides that are continuously updated. If a disk controller failure occurs at one site, the controller at the other site automatically takes over and accepts read/write I/Os. Global-active device assures that an up-to-date storage volume is always available and enables production workloads on both systems, while maintaining full data consistency and protection.

Hitachi Accelerated Flash

Hitachi Accelerated Flash features a flash module built specifically for enterprise-class workloads. Developed for Hitachi Virtual Storage Platform, Accelerated Flash is available for Hitachi Unified Storage VM and Hitachi Virtual Storage Platform family.

Accelerated Flash features innovative Hitachi-developed embedded flash memory controller technology. Hitachi flash acceleration software speeds I/O processing to increase flash device throughput.

Hitachi Storage Virtualization Operating System

Hitachi Storage Virtualization Operating System (SVOS) spans and integrates multiple platforms. It integrates storage system software to provide system element management and advanced storage system functions. Used across multiple platforms, Storage Virtualization Operating System includes storage virtualization, thin provisioning, storage service level controls, dynamic provisioning, and performance instrumentation.

Storage Virtualization Operating System includes standards-based management software on a Hitachi Command Suite (HCS) base. This provides storage configuration and control capabilities for you.

This solution uses Hitachi Dynamic Tiering, a part of Storage Virtualization Operating System, for data mobility. Separately licensed, Dynamic Tiering virtualizes and automates mobility between tiers for maximum performance and efficiency. Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Dynamic Tiering manages multiple storage tiers as a single entity. It presents a virtual volume with embedded smart tiering to monitor access and move data based on demand.

Dynamic Tiering automatically moves infrequently referenced data to lower cost tiers of storage. This data placement provides higher performance with lower costs. It also provides automatic wide-striping performance optimization.

When Tier 1 hardware uses a solid-state device or a flash module drive, active flash in Dynamic Tiering provides special care for write endurance. This complements using flash storage in your environment.

Storage Virtualization Operating System uses Hitachi Dynamic Provisioning (HDP) to provide wide striping and thin provisioning. Dynamic Provisioning provides one or more wide-striping pools across many RAID groups. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) without initially allocating any physical space. Deploying Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs).

Hitachi Thin Image

The high-speed, nondisruptive snapshot technology of Hitachi Thin Image snapshot software rapidly creates up to one million point-in-time copies of mission-critical information within any Hitachi storage system or virtualized storage pool without impacting host service or performance levels.

Because snapshots store only the changed data, the volume of storage capacity required for each snapshot copy is substantially smaller than the source volume. As a result, Thin Image snapshot software can provide significant savings over full volume cloning methods.

Thin Image snapshot copies are fully read/write compatible with other hosts. They can be used for system backups, application testing, and data mining applications while the business continues to run at full capacity.

Hitachi ShadowImage

High-speed, nondisruptive local mirroring technology of Hitachi ShadowImage heterogeneous replication, a part of Hitachi In-System Replication Bundle, rapidly creates multiple copies of mission-critical information within all Hitachi storage systems.

ShadowImage keeps data RAID-protected and fully recoverable, without affecting service or performance levels. You can split replicated data volumes from the host applications and used for system backups, application testing, and data mining applications, while business continues to run at full capacity.

Hitachi TrueCopy

Hitachi True copy remote replication, a part of Hitachi Remote Replication Bundle, is ideal for the most critical data situations when replication and backup of saved data are important. This software addresses these challenges with immediate real-time and robust replication capabilities.

This replicates between any storage systems within a virtualized storage pool managed by Hitachi Virtual Storage Platform family storage.

Hitachi Universal Replicator

Hitachi Universal Replicator is an advanced technology for asynchronously replicating data hosted on Hitachi Virtual Storage Platform family systems. Virtual Storage Platform manages the process of replicating the changes to the secondary site.

Universal Replicator software uses disk-based journaling and an optimized replication engine to reduce resource consumption and costs while increasing performance and operational resilience. The strengths of Hitachi Universal Replicator software are two key technical innovations: performance optimized, disk based journaling and a pull-style replication engine.

There is a potential for some data lag between remote and primary sites, particularly at longer distances. The recovery point objective is managed with the configuration of the data communication lines. When I/O activity at the primary site exceeds the capacity of the communication channel, the data is staged and moved to the secondary site in the same order as it was written at the primary site.

Hitachi NAS Platform

Hitachi NAS Platform is an advanced and integrated network attached storage (NAS) solution. It provides a powerful tool for file sharing, file server consolidation, data protection, and business-critical NAS workloads.

- Powerful hardware-accelerated file system with multi-protocol file services, dynamic provisioning, intelligent tiering, virtualization, and cloud infrastructure
- Seamless integration with Hitachi SAN storage, Hitachi Command Suite, and Hitachi Data Discovery Suite for advanced search and index
- Integration with Hitachi Content Platform for active archiving, regulatory compliance, and large object storage for cloud infrastructure

Hitachi Compute Blade 500

Hitachi Compute Blade 500 combines the high-end features with the high compute density and adaptable architecture you need to lower costs and protect investment. Safely mix a wide variety of application workloads on a highly reliable, scalable, and flexible platform. Add server management and system monitoring at no cost with Hitachi Compute Systems Manager, which can seamlessly integrate with Hitachi Command Suite in IT environments using Hitachi storage.

Hitachi Compute Blade 2500

Hitachi Compute Blade 2500 delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

Rack Optimized Server for Solutions, 2U Four Node

The rack optimized server for solutions, 2U four node, is an ultra-dense design equipped with four independent nodes. It creates the flexibility to set up different workloads independently in one 2U shared infrastructure, providing optimal data center performance per dollar.

It provides the following features:

- Up to four powerful dual Intel Xeon E5-2600 processor server nodes in a single 2U rack chassis
- Up to 24 hot-pluggable disks and flexible networking options
- Simple web interface for remote initialization, management and configuration

Rack Optimized Server for Solutions, 2U Single Node

The rack optimized server for solutions, 2U single node, is a rack mounted server designed for optimal performance and power efficiency. It supports up to 1.5 TB highly scalable memory capacity. It is powered by the Intel Xeon E5-2600 v3 processor product family for complex and demanding workloads. It supports flexible OCP and PCIe I/O expansion card options.

It provides the following features:

- Flexibility and performance drive solutions for many applications
- High-density shared power and cooling lowers energy costs while delivering optimal temperatures, even when a fan fails
- Redundant, hot-swappable drives and power supplies and dual-rotor fans keep non-stop business operations up and running

Hitachi Command Suite

Hitachi Command Suite manages virtualized storage and server infrastructures. With usability, workflow, performance, scalability, and private cloud enablement, Hitachi Command Suite lets you build sustainable infrastructures with leading storage technologies. It helps you flexibly align with changing business requirements and maximize return on IT investments.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Resources website.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Training and Certification website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 866 374 5822 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

© Hitachi Data Systems Corporation 2017. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. HDS, ShadowImage, and VSP are registered trademarks or trademarks of Hitachi Data Systems Corporation. Microsoft, Windows Server, and Office are registered trademarks or trademarks of Microsoft Corporation. All other trademarks, service marks, and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

AS-619-00, September 2017.